
Toward Virtual Community Knowledge Evolution

MICHAEL BIEBER, DOUGLAS ENGELBART, RICHARD FURUTA,
STARR ROXANNE HILTZ, JOHN NOLL, JENNIFER PREECE,
EDWARD A. STOHR, MURRAY TUROFF, AND
BARTEL VAN DE WALLE

MICHAEL BIEBER is an Associate Professor of Computer and Information Science at the New Jersey Institute of Technology (NJIT) in Newark. He conducts research in several related areas: hypermedia functionality, automatically generating hypermedia links and services for analytical applications and for digital libraries, Web engineering, incorporating hypermedia into World Wide Web applications, relationship analysis (as part of the software engineering process), supporting knowledge and learning within virtual communities, asynchronous learning networks and distance education, and infrastructures for future educational software. Dr. Bieber has published a number of papers in these and other areas. He is also active in several of these research communities, co-organizing conference minitracks and coediting special journal issues.

DOUGLAS ENGELBART is actively researching his lifelong goal of “augmenting the human intellect”—determining how technology can improve people’s effectiveness and increase the ability for people and organizations to build upon each other’s knowledge. His Bootstrap Institute is dedicated to helping organizations achieve this goal. He has received numerous awards, including the ACM Turing Award. Douglas Engelbart is well-known as the inventor of the computer mouse and the developer of one of the first hypermedia systems, NSL, in the early 1960s. He is one of the founders of the APRANET, and a pioneer in the areas of office automation and computer-supported cooperative work.

RICHARD FURUTA is a Professor in the Department of Computer Science, Director of the Hypermedia Research Laboratory, and Associate Director of the Center for the Study of Digital Libraries, all at Texas A&M University. He received a B.A. from Reed College in 1974, an M.S. in Computer Science from the University of Oregon in 1978, and a Ph.D. in Computer Science from the University of Washington in 1986. Dr. Furuta’s current areas of research include digital libraries and hypermedia systems and models.

STARR ROXANNE HILTZ is Distinguished Professor of Computer and Information Science, at NJIT, where she also directs the Ph.D. program in Information Systems. She does research on applications and social impacts of computer technology, publishing widely in journals including *Journal of Management Information Systems*, *MIS Quarterly*, *Communications of the ACM*, and *Management Science*. Her research interests currently include Group Support Systems, Asynchronous Learning Networks, and Pervasive Computing.

JOHN NOLL is Assistant Professor in the Computer Engineering department at Santa Clara University. He completed his Ph.D. in Computer Science at the University of Southern California, where he was a Research Associate with the ATRIUM laboratory at USC's Marshall School of Business. He has also taught at USC, UCLA, and the University of Colorado (Denver), and held positions at Perceptronics, Inc., Hewlett-Packard Laboratories, and Network Appliance, Inc. His research interests include software process, workflow, and computer-supported cooperative work.

JENNIFER PREECE is a Professor and Chair of Information Systems at University of Maryland, Baltimore County, where she researches and teaches classes in human-computer interaction and online communities. Her research focuses on the integration of usability and sociability in the design of online communities. Dr. Preece's current research projects include: assessing and supporting empathy in online patient support groups, design of tools for moderators, understanding and support lurking, and development of metrics to assess online activity. Her recent books include *Online Communities: Designing Usability, Supporting Sociability*, John Wiley & Sons, Chichester, England, and *Interaction Design: Beyond Human-Computer Interaction*, coauthored with Yvonne Rogers and Helen Sharp, due to be published by John Wiley & Sons in early 2002.

EDWARD A. STOHR holds a Bachelor of Civil Engineering degree from Melbourne University, Australia, and M.B.A. and Ph.D. in Information Science from the University of California, Berkeley. He is currently Associate Dean for Research and Academics at the Howe School of Technology Management, Stevens Institute of Technology, Hoboken, New Jersey. His research interests are centered on the problems of developing computer systems to support work and decision-making in organizations. He has been the editor of two books on decision support systems and has published articles in many leading journals. In 1992, Dr. Stohr served as chairman of the executive board of the International Conference on Information Systems (ICIS). He is on the editorial boards of a number of leading journals in the information systems field and has also acted as consultant to a number of major corporations.

MURRAY TUROFF is Distinguished Professor and Chair of the Information Systems Department at NJIT. He has been engaged in research and development of Computer-Mediated Communication Systems and related areas for the past three decades. Together with Starr Roxanne Hiltz, he is a coauthor of the award winning book, *The Network Nation* (MIT Press, 1993).

BARTEL VAN DE WALLE is Assistant Professor in the Information Systems Department of NJIT. His research interests are decision theory, group support systems, and the application of fuzzy set theory to negotiation in human and artificial multi-agent systems. His work is published in the *International Journal of Fuzzy Sets and Systems*, *Journal of Theory and Decision*, and the *Annals of Operations Research*.

ABSTRACT: This paper puts forth a vision and an architecture for a community knowledge evolution system. We propose augmenting a multimedia document repository (digital library) with innovative knowledge evolution support, including computer-mediated communications, community process support, decision support, advanced hypermedia features, and conceptual knowledge structures. These tools, and the techniques developed around them, would enable members of a virtual community to

learn from, contribute to, and collectively build upon the community's knowledge and improve many member tasks. The resulting Collaborative Knowledge Evolution Support System (CKESS) would provide an enhanced digital library infrastructure serving as an ever-evolving repository of the community's knowledge, which members would actively use in everyday tasks and regularly update.

KEY WORDS AND PHRASES: computer-mediated communication, concept maps, conceptual knowledge structures, decision analysis, digital libraries, hypermedia, hypertext, knowledge evolution, process modeling, professional societies, virtual communities, virtual educational communities, workflow.

DIGITAL LIBRARIES COMMONLY PROVIDE ACCESS to collections of published documents and multimedia items including digital representations of art, music, and so on [26, 27, 28]. Such depositories are usually "read only" in the sense that they are accessed and searched by users but developed and maintained by a professional group of librarians. In contrast, this paper describes a digital repository in which a community of users collaborate to share and evolve their knowledge in some domain of interest. The features of the enhanced digital library include computer-mediated communications, community process support, decision support, dynamic hypermedia, and conceptual knowledge structures. These tools, and the techniques we propose around them, will enable members of a virtual community to perform common community functions and to learn from, contribute to, and collectively build upon the community's knowledge [11].

A community's knowledge has both explicit and tacit components. The community's explicit knowledge includes its documents, recorded discussions, decision strategies, conceptual models, and defined workflows. Its implicit knowledge resides in the heads of the community members themselves but can be shared with others through processes of *socialization* (sharing experiences), *externalization* (articulating implicit knowledge into explicit concepts), *combination* (synthesizing and systematizing fragments of explicit knowledge), and *internalization* (turning explicit knowledge into tacit knowledge by applying it in real situations) [62]. Our proposed Collaborative Knowledge Evolution Support System (CKESS) is an enhanced digital library infrastructure supporting these four knowledge conversion modes through an ever-evolving repository of the community's knowledge, which members actively discuss and use in everyday tasks and regularly update. Thus, community members will be able to contribute to, discuss, and learn from the community's explicit knowledge, and share their implicit knowledge [66].

Virtuality has several definitions. For some, it indicates distance, requiring collaborators to communicate asynchronously ("different time, different place"). For some, it expresses the ability of computers to represent information in ways different from reality, with new tools allowing a broad range of different people to understand

complex or conceptual information and participate in exploring it [93]. For some, it indicates an organizational (or community) structure that is flexible enough to optimize individual and group performance under new and changing conditions [59, 71]. For some, it creates a sense of sharing experiences and perspectives, and emotional support between people working toward similar goals or solving problems together [33, 88]. In our view, participation in virtual communities should involve all these characteristics.¹

We broadly define a virtual community to include anyone actively interested in, or associated with, a group formed around a particular domain of interest. Dispersed or local, the community requires electronic support to implement a continuous meta-improvement strategy in its services. Thus we parallel Mowshowitz's view of virtual organizations—flexible organizations that actively seek flexible approaches to their own improvement [59].

In this paper, we consider two types of virtual communities: professional societies and virtual educational communities. A professional society is a special kind of virtual community, in which members participate to better understand its domain and improve the way they perform community-related tasks. The virtual community of a professional society may include nonmembers and organizations and often is many times larger than the professional society's membership. It can be thought of as an extended community of practice [100]. Virtual educational communities move beyond the common "push" mode of current distance learning systems to provide a truly interactive environment of mutual sharing and action learning. Professional societies and educational institutions have existing organizational infrastructures in place, and a mutuality of interests, which can help to maintain and support the CKESS.

CKESS should actively facilitate the growth of a virtual community into a Networked Improvement Community (NIC). Engelbart defines NICs as organizations that continuously improve the way they improve their products and services [24, 25]. Continuous meta-improvement occurs recursively at two levels. For example, in a professional society (1) members should find new ways to improve both the way they understand the domain and the way they perform their tasks, and (2) the professional society should continuously reevaluate the way it improves its community support. Meta-improvement arises from critically examining the processes underlying knowledge evolution and task support. CKESS will achieve continuous meta-improvement by forming augmented "meta-improvement discussions" aimed at improving or "evolving" these processes. Meta-improvement, thus, is realized through improving process innovation. We believe that this concept can bring about the widespread flexibility or "virtuality" Mowshowitz envisioned.

Research associated with the development and implementation of the proposed CKESS will shed light on seven major issues in the field of Digital Libraries and Virtual Communities:

1. The feasibility, effectiveness, and impact of supporting a virtual community with an enhanced digital library infrastructure—at the individual, group, and community level.

2. The evaluation of alternative approaches and guidelines that may be used by a virtual community to develop its membership, encourage participation, and effectively promote learning.
3. How an enhanced digital library infrastructure can support a professional society in educating its members, as well as nonmembers and students learning about its domain.
4. How an enhanced digital library-based meta-improvement infrastructure can transform a professional society into an NIC.
5. How multimedia repositories, full hypertext support, process support, decision support, computer-mediated communication, and conceptual knowledge structuring can be integrated into an enhanced, duplicable digital library infrastructure for supporting virtual communities.
6. The development of collaboration tools that allow a group to collectively understand its contributions to a common understanding, and to evaluate these contributions utilizing both appropriate social judgment scaling and voting tools.
7. The extension of concept mapping into full conceptual knowledge structures to facilitate collaborative knowledge evolution.

Most digital library efforts to date have concentrated on reengineering library features and procedures for electronic repositories. We believe that much more is needed. Our proposed CKESS, and the research associated with its development and deployment in existing virtual communities, will extend the scope of digital library support for collaborative knowledge evolution and continuous meta-improvement within virtual communities.

State of the Art

SEVERAL DOCUMENT REPOSITORIES (document bases with selective archival, search, and retrieval facilities) already exist in various communities. For some, such as the ACM Digital Library,² subscribers may not add materials except through a formal review process. Other efforts such as the Computing Research Repository (CoRR)³ [36] and NCSTRL⁴ [21] have been successful in encouraging community members to both add and use material, and the Open Archives Initiative (OAI)⁵ [47] is serving as the catalyst for development of standards for protocols that broaden these efforts. Our objective is to go even further and turn such repositories into full environments for knowledge evolution.

Many researchers have observed that knowledge management primarily is about people and cultural change rather than technical development [1, 99]. Research on online communities concerns itself explicitly with supporting people networking together to achieve a goal. Through this networking, knowledge is created and exchanged. Technology now plays an important role by supporting activities, recording knowledge, and developing organizational memory [2, 3, 90].

Key technical contributions to digital library efforts come from areas such as information retrieval/database management [67, 80], hypertext/hypermedia support [34, 64, 72], artificial intelligence [16] and information visualization [46, 83]. Although some digital libraries support collaborative use [58, 78, 87], most technological support concerns building collections rather than enhancing ways to use them. Studies exist concerning how a library should be used [48, 49] and how particular implementations are being used [12, 29, 37, 52], but to date, the library's role has been largely passive, rather than active, in directing the users' activities (see, however, some of the work on agency [74, 75]).

Much research is being done on synchronous environments such as MediaMOO,⁶ a professional online society for media researchers. Although the proposed CKESS will include a synchronous chat room, we believe it more important to incorporate asynchronous group support tools and computer-mediated communication (CMC). In this latter area, several issues have received scarce research attention: process support for virtual communities, the impact of digital libraries on virtual communities, and the idea of NICs.

A number of research projects investigate the provision of workflow and process support over the Internet via standard Web-based tools and protocols [57, 63, 76]. The intent is to support the activities of geographically dispersed "virtual corporations" and their customers [15, 103]. So far little work has been done to extend this concept to virtual communities. We propose that CKESSs use the capabilities of ad hoc workflow tools [41, 43, 50] to help coordinate the activities of the virtual community.

Since comprehensive deployment of digital libraries is only a recent phenomenon, few research efforts have studied the effects of their introduction on virtual communities. Studies to date have focused on the characteristics of physical communities [54, 55, 77, 102] or on evaluating specific digital library functions [12, 42, 45]. The closest research looks at the impacts of computer networking on scientific research communities [38, 89] and on the expected effect of digital libraries on the practices of scholarship [32].

Finally, no research has been conducted on the concept of NICs. Although much research and current business lore focuses on process improvement [19, 20], we have found no work besides Engelbart's [25] that studies systematic, continuous meta-improvement. The closest concepts are total quality management [35] and the SEI Software Capability Maturity Model [40]. The former lacks the emphasis of a continuous (and recursive) meta-improvement of an organization's improvement processes, whereas the latter is focused on traditional organizations, and thus far has not been applied to virtual communities of the kind CKESS supports.

Virtual Community Knowledge Evolution: The Vision

IN THIS SECTION, WE SET FORTH OUR VISION for virtual networked knowledge communities. The following hypothetical scenarios illustrate CKESS support for improvement in an educational community and professional society (scenarios 1–4) and

meta-improvement (scenario 5). The various tools mentioned in the scenarios will be described more fully in a later section.

Scenario 1: A Search Process for Digital Libraries

Mike Truss, the evening reference librarian at his university, was often asked how to find articles in the digital library for various homework projects. At his job now for over two years, Mike had long ago developed a mental checklist of questions that he asked every student in determining where to search in the digital library. Based on their responses, he knew which of the half dozen approaches each should take to find the information they needed, and he stepped the student through that procedure. When the university installed the CKESS digital library services package, Mike immediately designed an ad hoc workflow diagram expressing his mental checklist, as well as the half dozen search approaches. The workflow process included several decision points leading down different paths to the approaches. Based on a student's input, the process would automatically conduct part of the search, and guide the student on the other parts. Mike observed students using his workflow process, refining it as necessary. He was also able to refer students in the distance program to the process, which he found more effective than trying to guide them verbally over the phone. After a month, Mike formalized the process, set its access permissions to "general public," and posted a message describing it on the reference librarian's community discussion area. Within a few months, students from other areas of the country were using the process to locate articles more quickly in the digital library, without the help of reference librarians.

Scenario 2: Virtual Project Teams

Professor Smith teaches civil engineering in the distance learning program. Project teams in her bridge design course can rarely meet face-to-face as they live in different areas and have different work schedules. For her semester project she assigned teams of five students to thoroughly investigate a different kind of bridge. Each team was to prepare a Web site presenting its findings, as well as to contrast specifics of these findings with the reports of each other group. Virginia Hamlin led one of the project teams, which could never meet in person. Virginia designed a preliminary conceptual knowledge structure (concept map or CKS), sketching out the different aspects of a suspension bridge: its history, its structure, its physical characteristics, examples, designs, articles, research studies, and so on. She then drew up a preliminary ad hoc workflow process describing what each team member should do. The teammates followed her process, adding links and annotations to the various documents they found. Most of the documents were in their university's digital library. The team members constantly looked at each other's work, adding comments, interlinking related data and documents, and occasionally revising the CKS, which served as a visual overview of how all the different aspects of the project fit together.

Two-thirds of the way through the semester, Professor Smith gave all the teams read-only access to everyone else's project, so they could start interlinking related aspects of each. The project leaders began by studying each others' CKSs with the goal of constructing a common CKS. The students participated in a separate discussion area within the CKESS, in which a heated debate took place. Several alternate CKSs were proposed and discussed. After a week, Professor Smith directed the project leaders to vote on one CKS, that everyone then referred to as the "official overview" of the class as a whole. The official CKS provided an overview of bridge design and linked down to common aspects within each project group's CKS. Groups created links between their own sites and the official site, and added links directly among sites. Each link had a semantic type such as "predecessor," "alternate design," "contrasting view," "stress-test data," "scientific evidence," "underlying theory," "illustration," and so on.

Scenario 3: Developing a New Curriculum

Note: This scenario would apply for a curriculum committee within a single academic department or school, as well as for a curriculum committee within a professional society.

The International Engineering Professionals Association's Curriculum Committee decided that a new series of courses should be developed to train engineers in electronic commerce infrastructures. The committee comprised professionals and academics from across the country and therefore had to work asynchronously. Committee members began by proposing alternate CKSs illustrating the concepts they believed the curriculum should cover and how these concepts interrelated and built upon each other. They used CKESS's discussion environment to discuss aspects of each proposed structure, linking specific discussion comments to parts of specific CKSs. After several votes, they finally agreed on a single CKS for the new curriculum. Then the committee developed a series of courses and the general topics for each. Using the dynamic hypertext capabilities of CKESS, they linked in existing course syllabi and textbooks found in the community's digital library. They also linked in editorials and articles in the popular press and professional literature justifying certain aspects of particular courses and the curriculum as a whole. Academics on the committee found articles justifying other aspects in the educational research literature.

The committee then opened the curriculum up for discussion by the entire Association's membership. The CKS provided an overview of the curriculum's concepts. From the CKS, several hyperlinked guided tours emanated, providing alternate paths through selected lower-level documents within the proposal: one for professionals in the field, a second for professors, and a third for undergraduate students, who were also invited to review the proposal. Each guided tour linked together and annotated different sets of lower-level documents specifically for that group of stakeholders. Several online discussions ensued, and many annotations were made by the membership and responded to by the committee. The committee also made most of their original online discussion public so people could see the rationale and process

behind their proposal. Many comments were also made on specific discussion items. Based on this feedback, the committee made its final recommendation, which many universities went on to implement.

Scenario 4: Elections—Process Improvement

The International Engineering Professionals Association holds biannual elections according to their bylaws. A nominating committee creates a slate with two candidates per position. Community members return paper ballots to the Association headquarters. Elected to two-year terms, the election process is often new to officers. Luckily, CKESS includes several election processes, annotated by past officers. One describes the entire election process from start (forming a nominating committee and coordinating paperwork with the Association) to finish (turning over records to the new officers). From ongoing computer-mediated discussions, the community clearly understands the process better now. Before, it had mostly been a mystery because no one had easy access to it, which CKESS provides. Officers also state that the entire process, clearly recorded, is straightforward now, and that comments by past officers have provided a very useful “organizational memory” of potential problems. Community members also use CKESS to learn about the candidates, accessing both the research work and opinions written over the years in public conversations and newsletter articles. Candidates use CKESS to hold an asynchronous online debate on the issues, which many check at least once before voting. Some members add their own questions to the debate. One group even followed the CKESS-recorded process to place a third candidate on the online ballot for president. The membership voting in elections doubles.

Scenario 5: Meta-Improvement

The following scenario demonstrates meta-improvement because it implements a new process that evolved from a discussion on how to improve the way elections are run. The act of instituting such discussions was a new form of improvement for the organization, and these discussions then resulted in a variety of other improvements for the organization. Thus, the discussion capability is a meta-improvement resulting in a continuous stream of new improvements.

As a result of a “meta-improvement discussion” initiated by officers of the International Engineering Professionals Association, the executive board decided to revamp the nominating process for officer elections. The committee will begin soliciting nominations from the general membership through special computer-mediated discussions, allowing the community to express a level of preliminary support for candidates using computer-mediated voting tools. Any candidate with a minimal threshold of preliminary support will appear on the ballot.

In addition, the Association previously had two vice presidents—one for the Americas and one for the rest of the world. Based on a lengthy discussion over the past two years, which arose from representing two vice presidents on the election CKS, the

Association has instituted a third vice president. The second will now represent Europe. The third will represent the rest of the world. The election process (and CKS) has been rewritten accordingly, replacing the old versions in CKES. A lively computer-mediated discussion has formed around them, which for the most part is very positive.

Professional societies are examples of communities with a structure that has developed incrementally. Consequently, it is often difficult to change this structure (and many differing opinions will exist as to the most “logical” arrangement). The discussion and meta-improvement, such as that which our proposed architecture facilitates, should provide a means for community evolution.

Virtual Community Documents and Tasks

A VIRTUAL COMMUNITY’S DIGITAL LIBRARY REPOSITORY should give access to all of its community’s documents: books, journal articles, conference papers, audio tapes, videos, still pictures, course syllabi, and so on. But to support many of the everyday tasks of community members, as well as to form a community memory and knowledge base, the digital repository should be expanded to support computer-mediated communications, process, workflow and decision analysis capabilities, and conceptual knowledge structures.

Table 1 illustrates the range of tasks that are performed by individuals and the community as a whole within a professional society. Similarly, Table 2 illustrates the range of activities and tasks performed within virtual educational communities. Many tasks may be done by an individual, and all may be done by a formal or informal group. Each task can be done face-to-face or, given the proper group communication tools, by group members working at a distance. Each task can be done synchronously or, given the proper group communication tools, a virtual group can perform them in a coordinated fashion asynchronously, that is, working together, but at different times. Given the proper digital library technology, the materials used can reside at a distance from the person or people using them.

CKES’ workflow process support will allow community members to model, view, and to a limited degree, execute many of the tasks listed in Tables 1 and 2. In addition, CKES will provide processes to assist people in using the system itself, including the activities in Table 3.

Developing the Infrastructure

NO EXISTING APPROACH ADDRESSES THE FULL RANGE of knowledge support and evolution we envision. For example, most group support technologies focus on relatively simple forms of communication for conferencing, brainstorming, passing messages, and routing documents. The advanced structures needed for large dispersed groups to effectively work together on a variety of issues using a range of media forms do not exist [51, 95]. Research on digital libraries primarily concerns archiving and retriev-

Table 1. Individual and Community Tasks for a Professional Society

Individual tasks	Community tasks
Learning about the community's domain	Making a budget
Learning about the community's members	Conducting elections
Teaching a course in the community's domain	Creating the newsletter
Finding materials and appropriate references for writing a research paper on domain topics	Submitting materials to the newsletter
Determining the state of the art for topics by researchers or software developers	Proposing a task force or special project
Mentoring community members in writing papers that meet the community's norms	Running a conference or workshop
Looking for interesting research issues in this domain	Recruiting new society members

Table 2: Educational Community Tasks

Searching a digital library
Studying and learning about a subject
Designing and teaching a course
Developing and integrating course materials
Designing or doing homework projects
Holding class and recitation sessions
Authoring a research paper, article, or textbook
Developing a curriculum
Advising and mentoring
Discussing, commenting, and evaluating ideas/materials
Researching and forming concepts about a subject
Learning about and becoming a member of a community

Table 3. Digital Library Tasks

Adding documents and metadata to the digital library
Adding hypermedia features (comments, links, tours, and so on) to CKES
Conducting a discussion around a document, hypermedia feature, or concept map
Including CKES materials in a new or existing concept map
Formulating a decision analysis
Creating a new workflow process

ing documents, though some community services such as peer review and annotation are being researched [13, 44].

Little work has been done on workflow systems that are flexible enough to capture and support work processes in large communities outside formal organizations. Also, researchers are only now beginning to study the ad hoc workflows that could benefit informal groups [6, 105]. User interface design of broadly accessible systems still seems to shut out many potential members of society [70, 82]. Hypermedia systems

primarily concentrate on manually constructing links and other services among documents [10]. Similarly, few hypermedia systems provide linking for other kinds of organizational knowledge and none enable interlinking the range of knowledge components as we propose here.

We envision a pragmatic CKESS architecture to comprise a series of support tools that are independent of any particular digital library infrastructure. Integration will be ensured through a set of standard API (application programming interface) messages for the proposed CKESS modules. In the remainder of this section, we will show how tools from a number of different disciplines can contribute to the goals of the CKESS project.

To facilitate knowledge support and evolution within virtual communities, we propose an initial set of integrated tools, which includes (1) computer-mediated communication, (2) conceptual knowledge structures, (3) community process support, (4) decision analysis support, and (5) advanced hypermedia features. The integration of these components will represent a major advance [30, 65]. Whenever possible, we propose the use of existing technologies and systems and emerging WWW standards [56].

Computer-Mediated Communications and Collaborative Knowledge Structuring

Various computer-mediated communication structures and tools allow “virtual communities” to work together online, not only communicating about specific activities, but also building a form of collaborative knowledge base. Hiltz and Turoff [39] and Turoff [92] describe a number of fundamental structures and procedures “beyond e-mail and listservs,” which help groups organize and retrieve information while collaborating via computer networks for months or years. Such structures within NJIT’s ongoing Virtual Classroom™ project include segmentation into separate “conferences” organized by topic, with a branch or threading structure, which keeps all replies in a linked structure. Several field trials with scientific research communities show the ways these tools enabled groups of researchers and practitioners to enhance their productivity [38, 94].

The current generation of tools and structures break down, however, when large groups (such as, class sizes over 50) try to use them for extended and intensive information gathering and analysis. Thus, we propose to develop a set of richer and more powerful tools designed specifically to support professional communities of many hundreds to thousands of participants.

A centerpiece of the work is the development of CKS. CKS are concept maps that guide discourse. The CKS owner draws a semantic-structure depiction of the domain to be discussed, for example, using labeled circles representing the issues and subissues and labeled lines representing relationships among them. Everyone partaking in the discussion must enter their comments and replies under one of the categories in the CKS template (as with argumentation discourse systems such as gIBIS [18]). CKS will allow a community of knowledgeable individuals to individually contribute and collaboratively assess and evolve the knowledge that represents their common field

of endeavor. The CKS suite of tools will expose current issues or contrary hypotheses about relationships in the community's domain. (It is not always a feasible or desirable goal to achieve full consensus in large groups.)

CKS development will build upon computer-mediated communication, concept maps, semantic hypertext, argumentation, and structural modeling to support the following activities:

- Community members can propose and reach agreement on semantic structures [95] that categorize various portions of the knowledge domain (including communications, documents, and processes).
- Members can use these same templates as morphologies for the collection, organization, and retrieval of discussions and other elements within the digital library.
- Individuals considered to be “organizers of knowledge” can initiate, evolve, and modify a knowledge structure according to new collaborative understandings of the domain. Organizers can also initiate votes on relationships, and change content and relationships.
- Groups use voting and scaling tools to reach a mutual understanding and consensus on the semantic links and nodes that are incorporated in the templates.

CKESS will provide the guidance and interfaces necessary for any community member to “organize” a CKS. CKS will also enable different customized visualizations of the discussion structure based on the work, such as Card et al. [14], Sack [73], and Smith and Fiore [85], and scientific visualization in general. This is in-line with the definition of virtuality enabling new ways of viewing textual and non-textual information in order to bring sense to large amounts of information for a broad range of people.

These tools will not merely foster consensus, but expose current issues or contrary hypotheses about relationships in the community's domain. Relationships among elements in a domain will be multi-valued in that participants will be able to vote on both belief in their existence, as well as the relative strength or significance of the elements to other community members.

We strongly believe that CKSs incorporating an integrated collection of tools, presented in a well-designed interface, will allow a community of knowledgeable individuals to individually contribute and to collaboratively assess and evolve the knowledge that represents their common field of endeavor.

Workflow and Process Support

Process modeling and execution encompass a spectrum of process types, from simple repetitive processes addressed by conventional workflow management systems, to knowledge and communication intensive activities such as design, software development, and business decision-making [57, 76, 86, 101], and from “emergent” or ad hoc workflows to mission critical, structured, predefined processes [6, 79]. Also, workflows and processes do not exist in isolation, but rather require and produce

documents and other artifacts [61]. And, process descriptions alone do not capture the entire meaning of work activities, but often must be interpreted within the context of the documents they produce and consume [22].

The tasks shown in Tables 1 through 3 represent a similar spectrum, from relatively simple workflows such as adding documents to the knowledge base, to more complex processes such as managing the review process for a conference. In both of these examples, standard reusable processes to guide users, route communications, monitor progress, send reminders, and provide an audit trail, will be developed. A community will also engage in less-structured, more spontaneous, and human-intensive activities such as arranging meetings, collaborating to produce a progress report, and teaching a course in the community's domain. To support such activities, community members must be able to design their own processes using a graphical interface, automatically instantiate the processes, and then obtain machine assistance in executing instances of the process. Alternately, the constructed process models could be used to provide active guidance for members seeking to perform tasks off-line. Finally, relaxing constraints on user interactions even more, a workflow tool can be used to assist users without any predetermined rules or imposed structure. In this case, the logging capabilities of the system could be used to record usage patterns for later analysis, for example, using data mining techniques.

Several Web-based ad hoc workflow engines have been developed in research labs and commercially that could be used to support virtual communities [105]. Standards for workflow process invocation and control over the Internet have also been developed [104].

We propose enhancing CKES with process models describing how the tasks are performed, and providing active guidance through execution (enactment) of these models. Using techniques similar to those discussed in Scacchi and Noll [76], CKES community members will describe (model) processes for tasks such as those listed in Tables 1 and 2 in a high-level process modeling language. Process modeling languages allow users to specify the type and order of tasks that comprise the process, and the resources (documents, tools, and "products") that are required and produced by the process. Members can augment the models with narrative annotations, scripts, and links to related documents. This will enable the community to collect, document, and evolve best practices useful for improvement and training. An innovative feature of CKES' process support will be the integration of process enactment and process improvement. Because CKES should allow submission of new processes and improvements to existing processes, community members become both process performers and process engineers engaged in improving their own processes.

Decision Analysis Support

The list of community tasks tabulated in the fourth section involves decisions at various points during the execution of these activities. Daily experience has taught us that most such decisions are nontrivial. The individual decision-maker is usually confused by multiple and potentially conflicting decision criteria that prevent that person

from choosing a best decision alternative. Groups often fail to reach an agreed decision because the individual members' choices or rank orders of the decision alternatives vary considerably.

Individual and group decision problems, and models addressing these, have been studied in a variety of research contexts such as decision theory, preference modeling, human judgment, and organizational behavior. The problem of effective decision-making has long been recognized in business organizations, and various advanced information and decision support systems have been developed to address this issue. Today, we are witnessing the first decision-aiding experiments for virtual community members on the Web. Recommendation systems for shoppers at Amazon.com or CDNOW.com and a variety of other Web shops illustrate the interest in, and commercial value of, such systems. The enhanced digital library environment envisioned here involves decision analysis support for an actively collaborating, large community, extending the basic support available today.

Therefore, we propose enhancing CKESS with multi-criteria decision analysis (MCDA) models. MCDA models allow for the explicit recognition of multiple conflicting criteria, or at a group level, multiple conflicting group members. MCDA models have been widely applied at the individual and organizational level, and are well documented in theoretical and organizational literature [23, 84, 96, 97, 98].

Implementation of decision analysis support models will require two major steps. First, MCDA models must be developed as flexible tools or "components" that can be plugged into a group's workflow when a decision point is reached or when the need for more substantial decision support arises to any individual in the group. Second, these models must be connected to the discourse structures previously described, so that the communication preceding a certain decision is captured and the decisions taken are annotated. This connection between discussion and decision establishes an auditable "trail" of the decision process of the individual or the group. This in turn will greatly improve the understanding and the transparency of the decision process of a group and any of its members [91].

Dynamic Hypermedia Support

Hypermedia links represent relationships. Hypermedia navigation, structuring, and annotation features, such as user-declared linking, comments, trails, guided tours, and structural search [10, 60, 81] all provide ways of representing, conveying, and sharing knowledge.

The Dynamic Hypermedia Engine (DHE) automatically generates links and metadata, as well as other hypermedia features for analytical applications [7, 8, 31]. DHE executes separately from the applications it serves. A wrapper program for each application (such as, CKESS tool) integrates it into the DHE engine architecture. Applications or their wrappers then connect to DHE through servlets. DHE intercepts all screens and documents passing between the application and the user interface, and maps each appropriate element of the message to a hypermedia node or anchor. A Web browser wrapper integrates these anchors into the document being displayed. It

generates an HTML document, which it passes through the servlet to the user's Web browser. DHE automatically provides all hypermedia features to applications, which remain hypermedia-unaware and, in fact, often unchanged.

We will provide three types of hypermedia support for CKESS. First, DHE will automatically generate links within CKESS content as widely as possible—including the discussions, processes, conceptual knowledge structures, and digital library documents, which users create. This is done by writing a “wrapper” for each CKESS component enabling its integration with minimal modification. Second, as described in the following integration section, all CKESS objects can participate in, or be served by, hypermedia functionality. Third, additional hypermedia functionality, such as guided tours, annotations, and structural search, will be extended to all CKESS components, as well as to digital libraries in general [9].

Integration

Integrating all these components could follow a hypermedia modeling perspective. In the DHE described previously, for example, people can link nodes, make comments about nodes, place nodes on guided tours, display nodes within system overviews, and so on. Essentially, a node is an “element of interest.” In many traditional hypermedia systems, only documents would be considered nodes. For many user tasks, however—links, comments, guided tours, overviews, and so on—all could be considered elements of interest, and therefore people might want to recursively link them, comment on them, add them to a guided tour, or include them in an overview. To implement this, we either model any component as a subclass of a node or map it to a node.

We model integration within CKESS in a similar way. All objects within any CKESS subsystem will be a subclass of a node or map to a node. This includes:

- Repository: all multimedia documents, their components, attributes, access permissions, and so on;
- Conceptual knowledge structures: conceptual knowledge structures as a whole and component elements; participants and participant roles, and so on;
- Computer-mediated communications: discussion threads and items, participants, groups and participant roles, communication structures such as votes, and so on;
- Process and workflow: workflows, processes and their elements, resources, inputs and outputs, and so on; and
- Decision analysis: options, preferences, weights, and so on.

Hypermedia links will allow each of these components to involve the others. Thus, users will be able to start a computer-mediated discussion around any CKESS element or node, such as, portion of a document, a process, a process step, a hypermedia annotation or guided tour, or even part of another discussion. A CKS could contain any CKESS element. Similarly, users can comment on and include any CKESS element in a guided tour, and so on.

Clearly, certain elements appear across subsystems, such as participants. For any element not modeled in the initial digital library repository, we need to add a module

to support it in a coordinated manner. Also to achieve integration, we must ensure that every CKESS object—as well as any compound object containing it—(1) has a unique identifier across subsystems, (2) a recognized semantic type or set of semantic types, and (3) if the object is virtual or generated upon demand, there must be a general algorithm or set of commands for the DHE to regenerate instances of its semantic type or object class.

Deployment

THIS SECTION DESCRIBES AN APPROACH to implementing the CKESS architecture together with several important implementation issues. The idea is to start with a single, limited virtual community, build a prototype for that community, and then expand efforts to other communities. The precursor to any development, of course, is a detailed requirements analysis to ensure understanding of the target virtual community. Development should be conducted in a strong evaluative environment where one can use current user experiences and practices to aid in the evolution of the tools and the interface.

As CKESS is developed, it will be important to plan, conduct, and continually revise several alternative deployment plans with the cooperation of the target community. This involves several issues:

1. Announcing and advocating CKESS: How best does a community announce and advertise CKESS' goals, features, and services to its membership, especially those features with which they are unfamiliar, such as process support, decision analysis, and conceptual knowledge structures? How will one get a critical mass of community members to “buy into” the repository, add materials to it, and use it to access information through it whenever this would be useful to them?
2. Managing and maintaining CKESS: How best should a community manage the knowledge repository? This involves finding a group of people in the target community willing to install the system and be responsible for it. This group will need to make the deployment announcements and perform the advertising previously described. They will also need to set up a help service. Some communities may require moderators or facilitators [5, 17].

Guidelines are especially important as they provide a process for virtual communities to perform the above tasks. (To the extent possible, these processes should be formalized within CKESS itself, as Table 3 indicates.) We envision the need for the sets of guidelines shown in Table 4.

Evaluation

IN THIS SECTION, WE PRESENT A PRELIMINARY SET of propositions for CKESS research and an indication of methods for evaluating the extent to which these outcomes

Table 4. Necessary Sets of Guidelines for Supporting Virtual Communities

<ul style="list-style-type: none"> • Announcing and advocating an enhanced digital library to attract a critical mass of authoring and use • Populating an enhanced digital library with content • Managing an enhanced digital library • Using an enhanced digital library as a public relations tool for virtual communities 	<ul style="list-style-type: none"> • Providing a continuous meta-improvement structure for a professional society as an NIC • Supporting the tasks and processes of a professional society • Maintaining an enhanced digital library • Using an enhanced digital library for education
--	--

are attained within a virtual educational community. The goal of the proposed research program is to understand and design an enhanced digital library functionality that supports the particular needs of online “learning communities” and integrates with other common software tools, such as conferencing systems (also called discussion forums, or bulletin boards) for asynchronous interaction.

Table 5 lists an initial set of propositions and measures directed toward one or the other of our sample virtual communities.

By “critical mass,” we mean a substantial and relatively consistent user base [4, 53, 68, 69]. No consistent measure of critical mass exists in the literature. People generally agree that it reflects the number of participants or activity necessary for a community to function (or thrive). Target percentages of participation should be established as part of the requirements analysis.

It is necessary to measure both improvement and meta-improvement. Measuring improvement is one step beyond measuring usability of the system. Usability evaluates whether people find the system easy to use. Improvement evaluates whether people find that using CKESS results in better understanding of the domain, better work or a better research community. Measuring meta-improvement is one step beyond measuring improvement. Do people (or the community as a whole) find that CKESS’ ability to improve their understanding and tasks continuously improves? Stated recursively, will the target community over time, through CKESS, help people improve the way they improve their understanding, and improve the way they improve the way they perform tasks?

To explore how different groups of community members can judge improvement and meta-improvement, data could be collected from community leaders/officers (measured through periodic semi-structured interviews), general community members (measured through surveys), and if possible, people outside the community.

The research should apply both formative and summative evaluation techniques, including both qualitative and quantitative methods. Formative evaluation would be used to iteratively assess and improve the functionality and usability of the CKESS prototype. Summative evaluation would help assess usage, impacts, and satisfaction, in general, and the extent to which one achieves each of the goals, in particular.

Table 5. Initial Propositions and Measures

P1:	An enhanced digital library will improve satisfaction with online courses by students and faculty. <i>Measures: online surveys, faculty interviews.</i>
P2:	Enhanced digital library support will increase the use of such resources in online courses. <i>Measures: incorporation of digital library using assignments into syllabi of online courses; number of visitors to CKESS from participating courses.</i>
P3:	A critical mass of community members will add documents, links, and annotations, conversations, conceptual knowledge structures, and processes to an enhanced digital library in a sustained manner. <i>Measures: number and type of sources for objects added each month.</i>
P4:	The community's knowledge representations (that is, conceptual knowledge structures) and dialog will lead to constantly evolving knowledge accessed by the community. <i>Measures: number of concept maps and stored communications.</i>
P5:	A critical mass of community members will use an enhanced digital library consistently. <i>Measures: tracking repeat visits per month.</i>
P6:	A critical mass of community members will find the combination of hypermedia, computer-mediated communication, conceptual knowledge structures, and decision and process support in an enhanced digital library useful. <i>Measures: online survey, diaries.</i>
P7:	Members will perform community and individual processes more effectively when supported by ad hoc workflow automation tools. <i>Measures: user acceptance of workflow automation for routine processes, number of ad hoc processes generated, workflow engine log files, online survey, diaries.</i>
P8:	A virtual community can maintain its organizational memory effectively through an enhanced digital library that combines hypermedia, computer-mediated communication, conceptual knowledge structures, and decision and process support. <i>Measures: member/officer interviews, member survey.</i>
P9:	An enhanced digital library can support continuous meta-improvement of the virtual community. <i>Measures: member/officer interviews, member survey.</i>
P10:	An enhanced digital library will promote collaborative work among community members. <i>Measures: member survey, count of individual and collaboratively authored papers by community members over the period of the project, analysis of discussion participation, e-mail interviews with authors.</i>
P11:	An enhanced digital library will create new roles and forms of active participation in the community. <i>Measures: analysis of roles and artifacts created, numbers of sharable artifacts created.</i>
P12:	An enhanced digital library will improve the mastery of course materials by students. <i>Measures: comparative grade distributions, failure rates.</i>

Conclusion

THIS PAPER PROVIDES A VISION AND A POSSIBLE ARCHITECTURE for a community knowledge evolution system. Most digital library efforts to date have concentrated on reengineering library features and procedures for electronic repositories. Our proposed research will extend the scope of digital library support for collaborative knowledge evolution and continuous meta-improvement within virtual communities.

Although we focus here on educational communities and professional societies, our proposed knowledge support could certainly enhance other virtual communities

including medical and professional communities, and also organizations such as business and nonprofit organizations. All these communities and organizations have multimedia documents, processes, and discussions, and all potentially could take advantage of an enhanced digital library to achieve knowledge evolution and perhaps the status of becoming a networked improvement community.

Acknowledgments: The authors gratefully acknowledge partial funding support for this research by the Alfred P. Sloan Foundation, the NASA JOVE faculty fellowship program, the United Parcel Service, the New Jersey Center for Multimedia Research, the National Center for Transportation and Industrial Productivity at the New Jersey Institute of Technology, the New Jersey Department of Transportation, the New Jersey Commission of Science and Technology, and NSF grant 990925.

NOTES

1. Virtuality also carries the notion of 3-D immersion and simulation. We do not use this notion of virtuality here.
2. www.acm.org/dl.
3. xxx.lanl.gov/archive/cs/intro.html.
4. www.ncstrl.org or cs-tr.cs.cornell.edu.
5. www.openarchives.org.
6. www.cc.gatech.edu/fac/Amy.Bruckman/MediaMOO.

REFERENCES

1. Abell, A, and Oxbrow, N. People who make knowledge management work: CKO, CKT or KT. In J. Liebowitz (ed.), *Knowledge Management Handbook*. Boca Raton, FL: CRC Press, 1999, ch. 4.
2. Ackerman, M.S., and Halverson, C. Considering an organizations memory. In *Proceedings of the ACM 1998 Conference on Computer Supported Cooperative Work (CSCW'98)*. New York: ACM Press, 1998, pp. 39–48.
3. Ackerman, M.S., and Halverson, C. Reexamining organizational memory. *Communications of the ACM*, 43, 1 (2000), 58–64.
4. Ackerman, M.S., and Starr, B. Social activity indicators: Interface components for CSCW systems. In *Proceedings of the ACM Symposium on User Interface: Interface Components for CSCW Systems*. New York: ACM Press, 1995, pp. 159–168.
5. Berge, Z.L. The role of the moderator in a scholarly discussion group (SDG). Technical report, Berge Collins Associates, 1992, available at star.ucc.nau.edu/star.ucc.nau.edu/~mauri/moderate/zlbmod.html.
6. Bernstein, A. Executing programs with varying degrees of specificities: Populating the spectrum of specificity. Internal report. Center for Coordination Science, Massachusetts Institute of Technology, Cambridge, 1998.
7. Bhaumik, A.; Dixit, D.; Galnares, R.; Tzagarakis, M.; Vaitis, M.; Bieber, M.; Oria, V.; Krishna, A.; Lu, O.; Aljallad, F.; and Zhang, L. Integrating hypermedia functionality into database applications. In R.H. Sprague, Jr. (ed.), *Proceedings of the Thirty-Fourth Hawaii International Conference on System Sciences*. Los Alamitos, CA: IEEE Computer Society Press, 2001.
8. Bieber, M. Hypertext and web engineering. In K. Grønbaek, E. Mylonas, and F. Shipman III (eds.), *Proceedings of the Ninth ACM Conference on Hypertext and Hypermedia*. New York: ACM Press, 1998, pp. 277–278.
9. Bieber, M.; Galnares, R.; and Lu, Q. Service integration for virtual communities. In *Fourth International Workshop on Web Engineering, World Wide Web 10 Conference*, Hong Kong, May 2001.

10. Bieber, M.; Vitali, F.; Ashman, V.; Balasubramanian, H.; and Oinas-Kukkonen, H. Fourth generation hypermedia: Some missing links for the World Wide Web. *International Journal of Human Computer Studies*, 47, 1 (1997), 31–65.

11. Bieber, M.; Im, I.; Rice, R.; Goldman-Segall, R.; Paul, R.; Stohr, E.; Hiltz, S.R.; Preece, J.; and Turoff, M. Towards knowledge-sharing and learning in virtual professional communities. In R.H. Sprague Jr. (ed.), *Proceedings of the Thirty-Fifth Hawaii International Conference on System Sciences*. Los Alamitos, CA: IEEE Computer Society Press, 2002.

12. Bishop, A.P. Digital libraries and knowledge disaggregation: The use of journal article components. In I. Witten, R. Akscyn, and F. Shipman III (eds.), *Proceedings of the Third ACM Conference on Digital Libraries*. New York: ACM Press, 1998, pp. 29–39.

13. Borgman, C.L. *From Gutenberg to the Global Information Infrastructure: Access to Information in the Networked World*. Cambridge, MA: MIT Press, 1999.

14. Card, S.; Mackinlay, J.; and Shneiderman, B. *Readings in Information Visualization: Using Vision to Think*. San Francisco: Morgan Kaufmann, 1999.

15. Casati, F., and Discenza, A. Supporting workflow cooperation within and across organizations. In *Proceedings of the 2000 ACM Symposium on Applied Computing*, vol. 1. New York: ACM Press, 2000, pp. 196–202.

16. Chung, Y.-M.; Pottenger, W.; and Shatz, B. Automatic subject indexing using an associative neural network. In I. Witten, R. Akscyn, and F. Shipman III (eds.), *Proceedings of the Third ACM Conference on Digital Libraries*. New York: ACM Press, 1998, pp. 59–68.

17. Collins, M.P., and Berge, Z.L. Moderating online electronic discussion groups. Paper presented at the 1997 meeting of the American Educational Research Association (AREA), Chicago, March 1997, pp. 24–28.

18. Conklin, J., and Begeman, M. (1988). gIBIS: A hypertext tool for exploratory policy discussion. *ACM Transactions on Office Information Systems*, 6, 4 (1988), 303–331.

19. Davenport, T.H. *Process Innovation: Reengineering Work Through Information Technology*. Cambridge: Harvard University Press, 1993.

20. Davenport, T., and Beers, M. Managing information about processes. *Journal of Management Information Systems*, 12, 1 (Summer 1996), 57–80.

21. Davis, J.R., and Lagoze, C. NCSTRL: Design and deployment of a globally distributed digital library. *Journal of the American Society for Information Science*, 51, 3 (2000), 273–280.

22. Dourish, P.; Bentley, R.; Jones, R.; and MacLean, A. Getting some perspective: Using process descriptions to index document history. In *Proceedings of the International ACM SIGGROUP Conference on Supporting Group Work*. New York: ACM Press, 1999, pp. 375–384.

23. Dyer, J.; Fishburn, P.; Steuer, R.; Wallenius J.; and Zionts, S. Multiple criteria decision making, multiattribute utility theory: The next ten years. *Management Science*, 38, 5 (1992), 645–654.

24. Engelbart, D. Knowledge domain interoperability and an open hyperdocument system. In *Proceedings of the Conference on Computer-Supported Cooperative Work*. New York: ACM Press, 1990, pp. 143–156.

25. Engelbart, D. Toward high-performance organizations: A strategic role for groupware. In *Proceedings of Groupware '92*. San Mateo, CA: Morgan Kaufmann, 1992, pp. 77–100.

26. Fox, E., and Marchionini, G. (eds.). Toward a world wide digital library. *Communications of the ACM*, 41, 4 (1998), 28–32.

27. Fox, E., and Marchionini, G. (eds.). Digital libraries. *Communications of the ACM*, 44, 5 (2001), 30–32.

28. Fox, E.; Akscyn, R.; Furuta, R.; and Leggett, J. (eds.). Digital libraries. *Communications of the ACM*, 38, 4 (1995), 22–96.

29. Furuta, R.; Marshall, C.C.; Shipman, F.M., III; and Leggett, J.J. Physical objects in the digital library. In E. Fox and G. Marchionini (eds.), *Proceedings of the First ACM Conference on Digital Libraries*. New York: ACM Press, 1996, pp. 109–115.

30. Gaines, B.R.; Chen, L.-J.L.; and Shaw, M.L.G. Modeling the human factors of scholarly communities supported through the Internet and the World Wide Web. *Journal of the American Society for Information Science*, 48, 11 (1997), 987–1003.

31. Galnares, R. Augmenting applications with hypermedia functionality and meta-information. Ph.D. dissertation, New Jersey Institute of Technology, Newark, 2001.

32. Goh, D., and Leggett, J. Patron-augmented digital libraries. In *Proceedings of the Fifth ACM Conference on Digital Libraries*. New York: ACM Press, 2000, pp. 153–163.
33. Goldman-Segall, R. *Points of Viewing Children's Thinking: A Digital Ethnographer's Journey*. Mahwah, NJ: Lawrence Erlbaum, 1998.
34. Haas, S.W., and Grams, E.S. Page and link classifications: Connecting diverse resources. In I. Witten, R. Akscyn, and F. Shipman III (eds.), *Proceedings of the Third ACM Conference on Digital Libraries*. New York: ACM Press, 1998, pp. 99–107.
35. Hackman, R.J., and Wageman, R. Total quality management: Empirical conceptual and practical issues. *Administrative Science Quarterly*, 40, 2 (1995), 309–344.
36. Halpern, J.Y. CoRR: A computing research repository. *ACM Journal of Computer Documentation*, 24, 2 (2000), 41–48.
37. Hill, L.L.; Carver, L.; Dolin, R.; Frew, J.; Larsgaard, M.; Rae, M.A.; and Smith, T.R. Alexandria digital library: User evaluation studies and system design. *Journal of the American Society for Information Science (Special Issue on Digital Libraries)*, 51, 3 (2000), 246–259.
38. Hiltz, S.R. *Online Communities: A Case Study of the Office of the Future*. Human-Computer Interaction Series. Norwood, NJ: Ablex, 1984.
39. Hiltz, S.R., and Turoff, M. Structuring computer-mediated communication systems to avoid information overload. *CACM*, 28, 7 (July 1985), 682–689. [Reprinted in D. Marca and G. Bock (eds.), *Groupware: Software for Computer-Supported Cooperative Work*. Washington DC: IEEE Computer Society Press, 1992, pp. 384–393.]
40. Humphrey, W. Characterizing the software process. *IEEE Software*, 5, 2 (1988), 73–79.
41. Jablonski, S., and Bussler, C. *Workflow Management: Modeling Concepts, Architecture and Implementation*. London: International Thompson Press, 1996.
42. Jones, M.L.W.; Rieger, R.H.; Treadwell, P.; and Gay, G.K. Live from the stacks: User feedback on mobile computers and wireless tools for library patrons. In *Proceedings of the Fifth ACM Conference on Digital Libraries*. New York: ACM Press, 2000, pp. 95–102.
43. Jorgensen, H., and Carlsen, S. Emergent workflow: Planning and performance of process instances. In *Proceedings of Workflow Management*, Munster, Germany, 1999, available at www.wi.uni-muenster.de/is/tagung/workflow99/index.cfm/proceedings/wfm99proceedings.
44. Kling, R. Social relationships in electronic forums: Hangouts, saloons, work-places and communities. In R. Kling (ed.), *Computerization and Controversy: Value Conflicts and Social Choices*, 2d ed. San Diego: Academic Press, 1996, pp. 426–454.
45. Komlodi, A., and Marchionini, G. Key frame preview techniques for video browsing. In I. Witten, R. Akscyn, and F. Shipman III (eds.), *Proceedings of the Third ACM Conference on Digital Libraries*. New York: ACM Press, 1998, pp. 118–125.
46. Kumar, V.; Furuta, R.; and Allen, R.B. Metadata visualization for digital libraries: Interactive timeline editing and review. In I. Witten, R. Akscyn, and F. Shipman III (eds.), *Proceedings of the Third ACM Conference on Digital Libraries*. New York: ACM Press, 1998, pp. 126–133.
47. Lagoze, C., and Van de Sompel, H. The open archives initiative: Building a low-barrier interoperability framework. In *Proceedings of the First ACM/IEEE-CS Joint Conference on Digital Libraries*. New York: ACM Press, 2001, pp. 54–62.
48. Levy, D.M. I read the news today, oh boy: Reading and attention in digital libraries. In *Proceedings of the Second ACM Conference on Digital Libraries*. New York: ACM Press, 1997, pp. 202–211.
49. Levy, D.M. Heroic measures: Reflections on the possibility and purpose of digital preservation. In I. Witten, R. Akscyn, and F. Shipman III (eds.), *Proceedings of the Third ACM Conference on Digital Libraries*. New York: ACM Press, 1998, pp. 152–161.
50. Leymann, F., and Roller, R. *Production Workflow: Concepts and Techniques*. Upper Saddle River, NJ: Prentice Hall, 2000.
51. Majchrzak, A.; Rice, R.E.; Malhotra, A.; King, N.; and Ba, S. Technology adaptation: The case of a computer-supported inter-organizational virtual team. *MIS Quarterly*, 24, 4 (December 2000), 569–600.
52. Marchionini, G.; Nolet, V.; Williams, H.; Ding, W.; Beale, J., Jr.; Rose, A.; Gordon, A.; Enomoto, E.; and Harbinson, L. Content + connectivity => community: Digital resources for a learning community. In *Proceedings of the Second ACM Conference on Digital Libraries*. New York: ACM Press, 1997, pp. 212–220.
53. Markus, M.L. Toward a critical mass theory of interactive media: Universal access,

interdependence and diffusion. In J. Fulk and C. Steinfield (eds.), *Organizations and Communication Technology*. Newbury Park, CA: Sage, 1990, pp. 194–218.

54. Marshall, C.C. Making metadata: A study of metadata creation for a mixed physical-digital library. In I. Witten, R. Akscyn, and F. Shipman III (eds.), *Proceedings of the Third ACM Conference on Digital Libraries*. New York: ACM Press, 1998, pp. 162–171.

55. Marshall, C.C.; Price, M.N.; Golovchinsky, G.; and Schilit, B.N. Designing e-books for legal research. In *Proceedings of the First ACM/IEEE-CS Joint Conference on Digital Libraries*. New York: ACM Press, 2001, pp. 41–48.

56. McCray, A., and Gallagher, M. Principles for digital library development. *Communications of the ACM*, 44, 5 (2001), 49–54.

57. Miller, J.; Palaniswami, D.; Sheth, A.; Kochut, K.; and Singh, H. WebWork: METEOR2's web-based workflow management system. *Journal of Intelligent Systems*, 10, 2 (March–April 1998), 185–215.

58. Mooney, R., and Roy, L. Content-based book recommending using learning for text categorization. In *Proceedings of the ACM Digital Libraries 2000 Conference*. New York: ACM Press, 2000, pp. 195–204.

59. Mowshowitz, A. Virtual organizations: A vision of management in the information age. *The Information Society*, 10, 4 (1994), 267–288.

60. Nielsen, J. *Multimedia and Hypertext: The Internet and Beyond*. San Diego, CA: AP Professional, 1995.

61. Noll, J., and Scacchi, W. Specifying process-oriented hypertext for organizational computing. *Journal of Networking and Computer Applications*, 24, 1 (2001), 39–61.

62. Nonaka, I., and Konno, N. The concept of “Ba”: Building a foundation for knowledge creation. *California Management Review*, 40, 3 (Spring 1998), 40–54.

63. Object Management Group (OMG). jFLOW—Workflow Management Facility. OMG Business Object Domain Task Force, Object Management Group, Needham, MA, July 4, 1998.

64. Parunak, H.V.D. Don't link me in: Set based hypermedia for taxonomic reasoning. In P.D. Stotts and R.K. Furuta (eds.), *Hypertext '91 Proceedings*. New York: ACM Press, 1991, pp. 233–242.

65. Preece, J. Empathic communities: Balancing emotional and factual communication. Interacting with computer. *The Interdisciplinary Journal of Human-Computer Interaction*, 12, 1 (1999), 63–77.

66. Ram, S.; Park, J.; and Lee, D. Digital libraries for the next millennium: Challenges and research directions. *Information Systems Frontiers*, 1, 1 (1999), 75–94.

67. Ribeiro-Neto, B.A., and Barbosa, R.A. Query performance for tightly coupled distributed digital libraries. In I. Witten, R. Akscyn, and F. Shipman III (eds.), *Proceedings of the Third ACM Conference on Digital Libraries*. New York: ACM Press, 1998, pp. 182–190.

68. Rice, R.E. Communication networking in computer conferencing systems: A longitudinal study of group roles and system structure. In M. Burgoon (ed.), *Communication Yearbook*, vol. 6. Newbury Park, CA: Sage, 1982, pp. 925–944.

69. Rice, R.E. Computer-mediated communication system network data: Theoretical concerns and empirical examples. *International Journal of Man-Machine Studies*, 32, 6 (1990), 627–647.

70. Rice, R.E., and Katz, J. (eds.). *The Internet and Health Communication*. Thousand Oaks, CA: Sage, 2001.

71. Robey, D.; Boudreau, M.-C.; and Storey, V.C. Looking before we leap: Foundations for a research program on virtual organizations, electronic commerce. In G. St.-Amant and M. Amami (eds.), *Papers from the Third International Conference on the Management of Networked Organizations*. Atlanta: CIS Publications, Georgia State University, 1998, pp. 275–290.

72. Rutledge, L.; Ossenbruggen, J. van; Hardman, L.; and Bulterman, D.C.A. Practical application of existing hypermedia standards and tools. In I. Witten, R. Akscyn, and F. Shipman III (eds.), *Proceedings of the Third ACM Conference on Digital Libraries*. New York, ACM Press, 1998, pp. 191–199.

73. Sack, W. Discourse diagrams: Interface design for very large conversations. In R.H. Sprague Jr. (ed.), *Proceedings of the Thirty-Third Hawaii International Conference on System Sciences*. Los Alamitos, CA: IEEE Computer Society Press, 2000.

74. Sanchez, J.A.; Leggett, J.J.; and Schnase, J.L. AGS: Introducing agents as services provided by digital libraries. In *Proceedings of the Second ACM Conference on Digital Libraries*. New York: ACM Press, 1997, pp. 75–82.
75. Sanchez, J.A.; Lopez, C.A.; and Schnase, J.L. An agent-based approach to the construction of floristic digital libraries. In I. Witten, R. Akscyn, and F. Shipman III (eds.), *Proceedings of the Third ACM Conference on Digital Libraries*. New York: ACM Press, 1998, pp. 210–216.
76. Scacchi, W., and Noll, J. Process-driven intranets: Life cycle support for process reengineering. *IEEE Internet Computing*, 1, 5 (September 1997), 42–49.
77. Schiff, L.R.; Van House, N.A.; and Butler, M.H. Understanding complex information environments: A social analysis of watershed planning. In *Proceedings of the Second ACM International Conference on Digital Libraries*. New York: ACM Press: 1997, pp. 161–168.
78. Schnase, J.J.; Leggett, J.J.; Metcalfe, T.; Morin, N.R.; Cunnius, E.L.; Turner, J.S.; Furuta, R.K.; Ellis, L.; Pilant, M.S.; Ewing, R.E.; Hassan, S.W.; and Frisse, M.E. The CoLib project: Enabling digital botany for the twenty-first century. In J. Schnase, J. Leggett, R. Furuta, and T. Metcalfe (eds.), *Digital Library '94 Proceedings*. College Station, TX: Texas A&M University, 1994, pp. 108–118, available at www.csdl.tamu.edu/dl94/paper/colib.html.
79. Sheth, A.; Georgakopoulos, D.; Joosten, S.; Rusinkiewicz, M.; Scacchi, W.; Wiledon, J.; and Wolf, A. Report from the NSF Workshop on Workflow and Process Automation in Information Systems. *SIGMOD Record*, 25, 4 (December 1996), 55–67.
80. Shin, D.; Jang, H.; and Jin, H. BUS: An effective indexing and retrieval scheme in structured documents. In I. Witten, R. Akscyn, and F. Shipman III (eds.), *Proceedings of the Third ACM Conference on Digital Libraries*. New York: ACM Press, 1998, pp. 235–241.
81. Shipman, F.; Furuta, R.; Brenner, D.; Chung, C.; and Hsieh, H. Guided paths through web-based collections: Design, experiences, and adaptations. *Journal of the American Society of Information Sciences*, 51, 3 (2000), 260–272.
82. Shneiderman, B. Universal usability. *Communications of the ACM*, 43, 5 (2000), 84–91.
83. Shneiderman, B.; Feldman, D.; Rose, A.; and Grau, X. Visualizing digital library search results with categorical and hierarchical axes. In *Proceedings of the ACM Digital Libraries 2000 Conference*. New York: ACM Press, 2000, pp. 57–66.
84. Simpson, L. Do decision makers know what they prefer? MAVT and ELECTRE II. *Journal of the Operational Research Society*, 47, 7 (1996), 919–929.
85. Smith, M.A., and Fiore, A.T. Visualization components of persistent conversations. In *ACM CHI 2001 Proceedings*. New York: ACM Press, 2001, pp. 136–143.
86. Sommerville, I., and Rodden, T. Human, social and organizational influences on the software process. Technical report CSEG/2/1995, Cooperative Systems Engineering Group, Lancaster University, Lancaster, UK, 1995.
87. Sonnenwald, D.; Marchionini, G.; Wildemuth, B.; Dempsey, B.; Viles, C.; Tibbo, H.; and Smith, J.B. Collaboration services in a participatory digital library: An emerging design. In T. Aparac, T. Saracevic, P. Ingwersen, and P. Vakkari (eds.), *Proceedings of the COLIS 1999 Conference*. Lokve, Croatia: Benja Publishing, 1999.
88. Sproull, L., and Kiesler, S. *Connections: New Ways of Working in the Networked Organization*. Cambridge, MA: MIT Press, 1991.
89. Star, S.L., and Ruhleder, K. Steps toward an ecology of infrastructure: Design and access for large-scale collaborative systems. *Information Systems Research*, 7, 1 (March 1996), 111–138.
90. Stein, E.W., and Zwass, V. Actualizing organizational memory with information technology. *Information Systems Research*, 6, 2 (1995), 85–117.
91. Tsoukiàs, A. Preference modeling as a reasoning process: A new way to face uncertainty in multiple criteria decision support systems. *European Journal of Operational Research*, 55, 3 (1991), 309–318.
92. Turoff, M. Computer-mediated communication requirements for group support. *Journal of Organizational Computing*, 1, 1 (1991), 85–113.
93. Turoff, M. Virtuality. *Communications of the ACM*, 40, 9 (1997), 38–43.
94. Turoff, M.; Hiltz, S.R.; Bahgat, A.N.F.; and Rana, A. Distributed group support systems. *MIS Quarterly*, 17, 4 (December 1993), 399–417.
95. Turoff, M.; Hiltz, R.; Bieber, M.; Rana, A.; and Fjermestad, J. Collaborative discourse structures in computer mediated group communications. In R.H. Sprague Jr. (ed.), *Proceed-*

ings of the *Thirty-Second Hawaii International Conference on System Sciences*. Los Alamitos, CA: IEEE Computer Society Press, 1999.

96. Van de Walle, B., and Turoff, M. Towards group agreement: The role of preference analysis. In *Proceedings of the Eurofuse Workshop on Preference Modeling and Applications*, Grenada, Spain, 2001, pp. 85–92.

97. Van de Walle, B.; Heitsch, S.; and Faratin, P. Coping with one-to-many negotiations in electronic markets. In A.M. Tjoa and R.R. Wagner (eds.), *Proceedings of the Second Negotiations DEXIA Workshop*. Los Alamitos, CA: IEEE Computer Society Press, 2001, pp. 747–751.

98. Van de Walle, B.; Turoff, M.; Bieber, M.; and Rana, A. Integrating soft computing in group decision support systems: A taste of the challenges and opportunities. In *AAAI-97 Fall Symposium (MIT, Cambridge, MA)*. Boston: American Association for Artificial Intelligence Press, 1998, pp. 58–62. [Also available as AAI Technical Report TR FS-97-04.]

99. Walsh, J.P., and Ungson, G.R. Organizational memory. *Academy of Management Review*, 16, 1 (1991), 57–91.

100. Wenger, E. *Communities of Practice: The Social Fabric of a Learning Organization*. New York: Cambridge University Press, 1998.

101. Wolf, A., and Rosenblum, R. Process-centered environments (only) support environment-centered processes. In *Proceedings of the Eighth International Software Process Workshop*, Wadern, Germany, March 1993.

102. Wolfe, J.L. Effects of annotations on student readers and writers. In K. Anderson (ed.), *Proceedings of the Fifth ACM Conference on Digital Libraries*. New York: ACM Press, 2000, pp. 19–26.

103. Workflow Management Coalition (WfMC). Workflow and the Internet: Catalysts for radical change. White Paper, *Workflow Management Coalition*, Lighthouse Point, Florida, June 1998.

104. Workflow Management Coalition (WfMC). Workflow standard—Interoperability: WfXML binding. In L. Fischer (ed.), *The Workflow Handbook 2001*. Lighthouse Point, FL: Workflow Management Coalition, 2001.

105. Zhao, J.; Kumar, A.; and Stohr, E. A workflow-centric model of organizational knowledge distribution. In R.H. Sprague Jr. (ed.), *Proceedings of the Thirty-Third Hawaii International Conference on System Sciences*. Los Alamitos, CA: IEEE Computer Society Press, 2000.